

Article

A Multi-Objective Scoring Approach to Contract and Exposure-Aware Re-Ranking in Real-Estate Recommendation

Bogdan Arct ^{1,2} , Mateusz Bieniek ^{1,2} , Bartłomiej Kanabus ² , Aleksander Kozłowski ² , Piotr Wetmański ^{1,2} , Michał Kruk ¹ , Sylwia Stachowiak ³  and Jarosław Kurek ^{1,*} 

¹ Department of Artificial Intelligence, Institute of Information Technology, Warsaw University of Life Sciences, 02-776 Warsaw, Poland; michal_kruk@sggw.edu.pl (M.K.)

² Property Group sp. z o.o., ul. Adama Naruszewicza 27/101, 02-627 Warsaw, Poland

³ Department of Computer Science, University SWPS, ul. Chodakowska 19/31, 03-815 Warsaw, Poland

* Correspondence: jaroslaw_kurek@sggw.edu.pl

Abstract

Large online marketplaces increasingly rely on multi-stage ranking pipelines where a learned relevance model is complemented by business-aware constraints such as contractual pacing, exposure caps and commercial alignment objectives. This paper develops a second-stage, contract-aware re-ranking layer for real-estate recommendation that explicitly balances user–item relevance with plan fulfillment, lead value and operational guardrails. The proposed multi-objective re-ranker (PMOR) combines a calibrated base relevance score with multiplicative business adjustments and subtractive penalties for approaching contractual caps and for within-slate similarity. The method supports heterogeneous settlement models, including pay-per-action and fixed-fee contracts, via contract-specific weights. Because the scoring function is deterministic and structured, it admits exact component-wise contribution analysis and counterfactual ablations without relying on surrogate explainability methods. Offline evaluation on production logs from an anonymized marketplace covers 4219 recommendation requests and 184,147 candidate items, joined with daily business snapshots using an as-of strategy to prevent look-ahead bias. Under a profit proxy based on effective lead value, position discounting and billability, PMOR achieves an indexed expected-revenue proxy of 487.4 (baseline = 100), corresponding to a lift of 387.4% over a model-only baseline and 48.4% over a legacy production re-ranker (LPR). The gain is primarily associated with improved billable exposure, increasing the share of billable positions in TOP-3 to 78.22% compared with 37.85% for LPR. We discuss parameter sensitivity, operational considerations and limitations of offline proxy objectives for deployment.

Keywords: recommender systems; re-ranking; multi-objective ranking; exposure control; pacing; diversity; contract-aware optimization; marketplace commercial alignment



Academic Editors: Costas Vassilakis and Dionisis Margaritis

Received: 4 June 2026

Revised: 5 July 2026

Accepted: 8 July 2026

Published: 11 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Modern online marketplaces rarely optimize recommendation lists solely for predicted user preference. Real-world deployments are subject to contract terms, budget constraints and operational objectives that a single relevance model cannot capture. In two-sided environments, ranking decisions mediate user utility, provider outcomes and commercial alignment, which motivates explicit business-aware objectives and guardrails [1–5]. In real-estate platforms, these pressures are particularly pronounced because inventory is diverse

across location, property type, price segment and developer, and the decision process is information-intensive; surveys and recent system papers highlight the breadth of modeling choices and the importance of how homebuyers search and consume information [6–10].

A common engineering pattern for meeting such requirements is a multi-stage ranking pipeline. A first stage retrieves candidates and computes a relevance score, while later stages implement pre-ranking and re-ranking logic under latency constraints and operational state [11–19]. Industrial datasets and end-to-end frameworks further emphasize that realistic evaluation must account for the full pipeline from retrieval to serving [20]. Because platforms often rely on logged interactions to improve ranking policies, learning and evaluation with logged (bandit) feedback has also been studied [21] and neural re-ranking has been investigated across multiple application domains [22].

Parallel to pipeline design, advances in base recommendation models—including attention-based similarity formulations [23], task-oriented graph embeddings and multi-relational graph neural networks [24,25] and more recent LLM-enhanced graph recommendation frameworks [26–28]—aim to improve relevance and representation quality. Meta-learning and rapid adaptation strategies explored in neighboring domains suggest additional avenues for coping with market heterogeneity and distribution shift [29]. Regardless of the base model family, however, real deployments still require a controllable second-stage mechanism for enforcing contractual constraints, exposure caps and pacing.

The ranking literature provides foundations for relevance optimization via learning-to-rank approaches, including early neural methods and listwise objectives [30,31]. In practice, objectives are multi-dimensional and often require explicit multi-objective optimization or scalarization strategies [32,33]. Recent work explores multi-objective and controllable re-ranking policies (e.g., rejection-sampling, Pareto-based formulations, policy hypernetworks and meta-learned ranking) [34–37] and slate-aware decision formulations [38]. A complementary line of work addresses redundancy by diversifying ranked outputs using heuristic and probabilistic mechanisms [39–41].

Exposure is also an economic resource in two-sided platforms. Constraints on attention allocation have been studied under fair exposure and equity of attention and extended to personalized or provider-aware settings [42–52]. Longitudinal analyses further suggest that exposure bias can compound over time, motivating proactive control mechanisms [53]. From an operational perspective, low-latency constrained ranking and dedicated solvers aim to make such constraints compatible with real-time serving [54]. Related work in advertising explores adaptive exposure control at scale, highlighting additional design patterns for real-time allocation [55].

Budget and pacing constraints have been explored extensively in auction markets and sponsored search, where delivery must be smoothed under uncertainty and budgets must be respected [56–63]. Although real-estate recommendation is not an auction, similar pacing logic is relevant when lead caps, plan fulfillment and contractual settlements interact with daily traffic patterns.

Finally, the notion of “contracts” in real estate is increasingly digital. Blockchain and smart-contract approaches have been surveyed and prototyped for property transfer, escrow and ownership verification, alongside broader discussions of tokenization [64–74]. This direction raises security, privacy and specification challenges for smart and legal contracts [75–78]. Our work is orthogonal to blockchain execution: we focus on the ranking-time policy layer that allocates exposure under conventional commercial contracts and operational caps.

This study formalizes and evaluates a business-aware re-ranking layer for a real-estate marketplace recommender operated by an anonymized industry partner. The central design goal is to provide a ranking mechanism that is (a) contract-aware, handling heterogeneous

settlement models; (b) operationally robust, preventing “over-concentration” on a single investment or developer by enforcing exposure caps and pacing; (c) diversity-aware, discouraging near-duplicate slates that degrade user experience; and (d) interpretable in the sense that each score component has a clear meaning and can be attributed exactly to a module of the decision rule.

The work contributes in three complementary directions. First, it proposes a parameterized multi-objective scoring function that combines calibrated relevance with business drivers and safety penalties, with explicit handling of pay-per-action versus fixed-fee settlement. Second, it provides an online sequential slate construction procedure that implements diversity and cap constraints at inference time, without requiring model retraining when business policy shifts. Third, it introduces a deterministic component-wise contribution analysis and counterfactual ablations that quantify how each module of the score affects ranking decisions, enabling targeted tuning and risk assessment.

Offline results show that the proposed re-ranker (PMOR) substantially improves a business-oriented proxy metric compared with both a model-only baseline and a legacy production re-ranker (LPR), largely by increasing the fraction of recommendations directed to billable inventory. The remainder covers the problem setting, method, evaluation protocol, results and deployment implications.

2. Related Work

Business-aware re-ranking sits at the intersection of learning-to-rank, recommender systems, constrained optimization and exposure control. The literature spans model-based ranking, listwise optimization, diversification and fairness/exposure constraints, as well as commercial-alignment objectives typical in advertising and marketplaces.

2.1. Multi-Stage Ranking in Recommender Systems

Large-scale recommender systems often use a candidate-generation stage followed by a ranking or re-ranking stage to meet latency constraints and incorporate additional features available only at later stages. Industrial systems commonly adopt this architecture to separate representation learning from policy logic and to enable frequent updates to business rules without retraining deep models [12,13]. Two-stage approaches also support efficient retrieval from large catalogs, followed by more expensive scoring on a narrower candidate set [11].

In the context of slate recommendation, decision-making operates on ordered lists rather than independent items. Listwise or slate-aware objectives aim to optimize the utility of the entire list under position bias and interactions among items [31,38]. However, direct slate optimization can be difficult to implement when the objective must incorporate real-time constraints such as contract caps and pacing, which depend on operational state and can be discontinuous.

2.2. Learning to Rank and Multi-Objective Ranking

Learning-to-rank has developed a rich family of pairwise and listwise methods, including RankNet and LambdaRank, that optimize ranking metrics such as NDCG [30]. In many practical settings, the target objective is multi-dimensional and involves trade-offs between relevance and business outcomes. Scalarization, constrained optimization and multi-objective formulations are standard tools to combine objectives, although they can be sensitive to scaling and require careful calibration [32]. In recommender systems, objectives may include not only accuracy but also coverage, novelty, diversity and business measures such as value or operational outcomes [1].

Our approach uses an explicit parametric scalarization at the re-ranking stage. The design emphasizes interpretability and operational controllability rather than end-to-end learned weights, reflecting the need to adjust strategy at high frequency and to enforce constraints that are not easily learned from historical data alone.

2.3. Diversity and Redundancy Control

Diversification aims to reduce redundancy and increase perceived quality by presenting varied items. Classical methods include maximal marginal relevance (MMR), which trades off relevance against similarity to already-selected items [39]. Diversification in recommender systems has been studied extensively, including topic diversification and hybrid relevance-diversity ranking [40]. Probabilistic models such as determinantal point processes provide a principled way to encourage diversity, although their computational cost can be nontrivial for real-time ranking [41].

In real-estate recommendation, redundancy often arises from multiple listings within the same investment or highly similar geographic regions. The diversity mechanism in PMOR is implemented as a discrete similarity penalty during sequential slate construction, enabling direct enforcement of “no near-duplicate” behavior in TOP-K.

2.4. Exposure Constraints, Fairness and Marketplace Dynamics

Exposure is not only a user-experience issue but also an economic resource distributed among suppliers. Several lines of work study fair exposure and attention allocation in ranked lists, including constraints that equalize exposure according to relevance or protected attributes [42,43]. In marketplaces, fairness constraints can also be framed as supplier fairness, requiring that smaller providers receive a minimum level of exposure or that exposure follows contractual commitments.

Our setting is contract-driven rather than demographically-driven fairness. The method incorporates normalized business signals as explicit features that influence exposure distribution. The goal is not to define a normative fairness principle but to implement a tunable policy that aligns exposure with operational targets while maintaining user relevance.

2.5. Commercial-Alignment Ranking and Pacing

Commercial-alignment ranking often appears in advertising systems and sponsored search, where expected revenue is a function of value per click (or per conversion), predicted conversion probability and position bias. Classic formulations involve expected value under position discounts, frequently expressed via DCG-style weights or click models [21]. Budget pacing controls the rate of delivering impressions or opportunities over time to avoid depleting budgets early in a period and to maintain smooth delivery under uncertain demand [56]. In marketplace recommendation, pacing can be interpreted similarly as a soft regulator of exposure relative to planned targets.

This work draws on these principles but adapts them to a recommendation context with heterogeneous settlement models and with explicit cap and diversity penalties that can strongly adjust the slate. The offline evaluation uses an expected-revenue proxy that accounts for effective lead value, estimated conversion probability, position discounting and billability under cap constraints.

2.6. Synthesis and Research Gap

Taken together, the reviewed literature addresses the ingredients of business-aware ranking in isolation. Multi-stage and learning-to-rank work optimizes relevance and list utility but does not natively encode contractual settlement or pacing state; multi-objective and Pareto-based re-rankers expose trade-offs but are typically tuned for offline

objectives and are expensive to steer at serving time; diversity and fair-exposure methods control redundancy and attention allocation but treat exposure as a fairness rather than a contractual resource; and pacing and budget work from auctions assumes an auction mechanism that real-estate recommendation does not have. The main limitation of prior work, for our setting, is therefore the lack of a single, controllable second-stage layer that jointly enforces heterogeneous contract settlement, exposure caps, pacing and diversity while remaining exactly interpretable.

Recent multi-objective recommenders built on graph neural and graph-attention architectures, such as [79], likewise optimize several objectives jointly, but they embed the trade-off inside the relevance model at training time; PMOR instead applies its multi-objective trade-off as a deterministic, model-agnostic second stage, so commercial policy can change without retraining the base model.

PMOR targets this gap. Its novelty is not a new relevance model but the combination of (i) contract-aware scalarization that handles pay-per-action and fixed-fee settlement in one scoring rule, (ii) inference-time slate construction that enforces caps and diversity without retraining, and (iii) deterministic, component-wise attribution of every ranking decision. Relative to learned or population-based multi-objective re-rankers, PMOR trades some frontier optimality for operational controllability and exact explainability, which is the property the deployment setting requires.

3. Materials and Methods

This section is organized as follows. We first fix notation and the problem setting (Section 3.1), then describe the offline sample, the business-aware multi-objective scoring function and its contract-aware weights, the feature and guardrail construction, and the sequential slate-construction procedure that produces the final TOP-K list. The reference-ranker and parameter-selection subsections define the evaluation setup rather than the PMOR method itself; the offline profit proxy that serves as the evaluation objective is now defined together with the research questions in the validation section (Section 4).

Figure 1 gives a high-level view of the pipeline. PMOR is a deterministic second-stage layer: for each candidate it starts from a calibrated base relevance score, applies multiplicative business adjustments (plan-gap, contract value, profitability) and subtractive guardrail penalties (cap and within-slate diversity), and then builds the slate greedily one position at a time so that diversity and cap constraints are enforced against the partial slate.




Figure 1. Overview of the PMOR re-ranking pipeline: a calibrated base relevance score is combined with multiplicative business adjustments and subtractive guardrail penalties into an item score, which drives greedy sequential slate construction under cap and diversity constraints.

The remainder of this section follows the pipeline of Figure 1 step by step. Equation (1) states the assembled item score (the output of Step 4 below); its inputs are then detailed in pipeline order: (i) the calibrated base relevance $q(u, i)$ (Step 1) in Section 3.4.1; (ii) the multiplicative business adjustments—plan-gap, contract value and profitability (Step 2)—in Sections 3.4.2 and 3.4.3, together with the contract-aware weights in Section 3.3; (iii) the subtractive guardrail penalties for cap proximity and within-slate diversity (Step 3) in

Sections 3.4.4 and 3.4.5; (iv) the assembled item score $s(u, i)$ (Step 4), Equation (1) in Section 3.3; and (v) the greedy sequential slate construction with pacing ρ (Step 5) in Section 3.5. Each subsection below is annotated with the pipeline step it implements.

3.1. Problem Setting and Notation

Consider a user request u that triggers a recommendation call. The system retrieves a candidate set $C_u = \{i_1, \dots, i_n\}$, where each candidate i corresponds to a property-level listing. Each listing belongs to a higher-level business entity such as an investment or campaign. We denote that entity by $\text{inv}(i)$ and similarly allow grouping by developer $\text{dev}(i)$ and city $\text{city}(i)$ when such identifiers exist.

A base recommender produces a relevance score $q(u, i)$ for each candidate. The role of the re-ranking layer is to produce an ordered slate $L_u = (i_{(1)}, \dots, i_{(K)})$ of length K that balances relevance with business goals and operational constraints. In our experiments, $K = 3$ is the primary reporting cutoff because TOP-3 dominates exposure in the deployed interface, though the method is not specific to $K = 3$.

We use TOP-3 because the deployed interface concentrates almost all user attention on the first three positions, so it is the cutoff where re-ranking decisions have the largest exposure impact; larger K dilutes the measured effect without changing the scoring rule, and the method applies unchanged to other slate sizes.

We also use $L \subseteq L_u$ to denote the partial slate currently under construction during sequential selection; $L = \emptyset$ at the start of a request and grows one item at a time as Algorithm 1 proceeds. The diversity penalty $D(i, L)$ is evaluated with respect to this partial slate, so it is implicitly request-scoped: every request u begins with $L = \emptyset$ and constructs its own slate independently of other requests.

Algorithm 1 Sequential slate construction for the proposed multi-objective re-ranker (PMOR)

Require: Candidate set $C_u = \{i\}$ with features (q, r, g, m, v, c) , cap state for computing $C(i)$, parameters Θ , slate length K

```

1:  $L \leftarrow []$ 
2: while  $|L| < K$  and  $C_u \neq \emptyset$  do
3:   for all  $i \in C_u$  do
4:     compute  $C(i)$  from  $\text{cap\_ratio}(\text{inv}(i))$  via Equation (11)
5:     compute  $D(i, L)$  via Equation (13)
6:     select  $(\beta, \eta)$  using Equation (2) based on  $c(i)$ 
7:      $S(u, i) \leftarrow \text{Equation (1)}$ 
8:   end for
9:    $i^* \leftarrow \arg \max_{i \in C_u} S(u, i)$ 
10:  append  $i^*$  to  $L$ ; remove  $i^*$  from  $C_u$ 
11:  update counters used by cap state (if online state is maintained)
12: end while
13: return  $L$ 

```

The business state includes, for each investment or billing entity, contract type, budget and plan information and progress toward caps. Let $c(i) \in \{\text{cpa}, \text{fixed}\}$ denote the settlement model associated with the investment of item i , where cpa denotes pay-per-action contracts and fixed denotes fixed-fee or subscription-like contracts. Let $g(i) \in [0, 1]$ represent a normalized plan-gap signal, higher when plan fulfillment is behind target. Let $r(i) \in [0, 1]$ represent a normalized profitability or lead-value signal. Additional features include a bounded contract-value signal $m(i) \in [-1, 1]$ and a premium flag $v(i) \in \{0, 1\}$.

Operational guardrails include an exposure cap penalty $C(i) \in [0, 1]$ derived from cap progress and a within-slate diversity penalty $D(i, L) \in [0, 1]$ derived from similarity to

already-selected items. A pacing signal $p(h) \in [-1, 1]$ is available at time h (e.g., hour-of-day) as a soft regulator of delivery relative to planned exposure.

Table 1 collects the symbols used throughout the paper for quick reference.

Table 1. Notation used throughout the paper.

Symbol	Meaning
u	user request
i	candidate item (property-level listing)
$\mathcal{C}_u$	candidate set for request u
$\text{inv}(i)$	investment containing item i
$\text{dev}(i)$	developer associated with item i
$\text{city}(i)$	city identifier for item i
$q(u, i)$	calibrated base relevance score, $q \in [0, 1]$
$S(u, i)$	PMOR score (Equation (1))
L_u	final slate of length K returned for request u
L	partial slate under sequential construction, $L \subseteq L_u$
K	slate length ($K = 3$ in main evaluation)
$c(i)$	settlement model, $c(i) \in \{\text{cpa}, \text{fixed}\}$
$g(i)$	plan-gap signal, $g(i) \in [0, 1]$
$r(i)$	profitability/lead-value signal, $r(i) \in [0, 1]$
$m(i)$	contract-value signal, $m(i) \in [-1, 1]$
$v(i)$	premium flag, $v(i) \in \{0, 1\}$
$C(i)$	cap-proximity penalty, $C(i) \in [0, 1]$
$D(i, L)$	within-slate diversity penalty, $D(i, L) \in [0, 1]$
$p(h)$	pacing signal at hour h , $p(h) \in [-1, 1]$
$\pi(i)$	effective lead price
$b(i)$	billability indicator, $b(i) \in \{0, 1\}$
w_j	DCG position weight at rank j
$\hat{R}_A(u)$	per-request profit proxy under ranker A (Equation (25))
$\hat{R}_A$	dataset-level profit proxy under ranker A (Equation (26))

3.2. Data Sources and Offline Sample Construction

The study uses anonymized production logs from a large online real-estate marketplace. Logs are organized at the level of recommendation requests. Each request yields a list of candidate properties along with base scores and an existing production score that reflects legacy business logic. To incorporate contract state and progress toward plan, logs are joined with daily business snapshots containing budget, lead price, lead limits and realized leads-to-date. The join is performed using an as-of backward strategy, matching each request date to the most recent snapshot not after the request timestamp. This strategy reduces the risk of look-ahead bias because future contract updates are not used when evaluating past requests.

Two practical issues motivate explicit preprocessing. First, base model scores may contain a sentinel value indicating missing or invalid output. Second, some business fields are missing for a subset of investments on certain days due to data sparsity or delayed updates. The pipeline therefore applies consistent fallback rules and clipping to maintain bounded feature ranges.

3.2.1. Offline Sample Summary

Table 2 reports key statistics for the offline sample used in the main evaluation. The sample consists of one request per user identifier, reflecting a common “latest request” sampling approach that reduces correlation among requests and simplifies per-request comparison.

Table 2. Summary of the offline evaluation sample constructed from production logs and daily snapshots.

Statistic	Value	Notes
Total candidate rows	184,147	Rows correspond to property candidates across all requests.
Unique users	4219	One request per user in the sampled dataset.
Unique requests	4219	Request-level unit for evaluation and pairwise comparison.
Unique investments	1388	Identified via investment identifiers in logs.
Average candidates per request	43.66	Computed as 184,147/4219.
TOP-K for evaluation	3	Main reporting cutoff in the production interface.
Snapshot coverage	49 days	Daily snapshots span a contiguous date range.
Unique offers in snapshots	1859	Offer-level identifiers in business snapshots.
Settlement mix in snapshots	51.03% CPA 48.97% fixed-fee	Measured on snapshot rows, not on exposures.

3.2.2. Special Values and Missingness

Base relevance scores may contain a sentinel value -1 indicating that the model score is unavailable. In the offline logs, this sentinel appears in 27.24% of rows and must be handled to avoid undefined behavior under exponentiation. The evaluation maps sentinel values to 0 and subsequently clips all relevance and business scores to $[0, 1]$ as needed.

Table 3 summarizes the prevalence of sentinel and boundary values in key score columns. Boundary values are informative because several pipeline components are explicitly normalized to $[0, 1]$ and may saturate at 0 or 1 under particular business states.

Table 3. Counts of sentinel and boundary values in the offline log scores.

Column	-1 Count	0 Count	1 Count	Non-Null
Base relevance score (raw)	50,156	24	212	184,147
Legacy production score	0	6	2	184,147
Plan-gap score	0	10,089	18,976	184,147
Profitability score	0	15,524	3262	184,147

3.3. Business-Aware Multi-Objective Scoring Function

The proposed re-ranking policy is expressed as a scoring function that transforms base relevance into a final decision score, incorporating business drivers and guardrails. The design follows three principles. First, relevance should remain a central component of the objective, balanced with business constraints when those must be enforced. Second, major business drivers should be contract-aware and bounded to facilitate stable parameter tuning. Third, hard-to-reason-about exposure dynamics should be controlled via explicit penalties rather than left implicit in a learned model.

3.3.1. Score Definition

This subsection specifies Step 4 of the pipeline (Figure 1): the assembled item score, whose inputs (Steps 1–3) are detailed in the subsequent subsections.

For a request u and a candidate item i , the PMOR score is defined as:

$$S(u, i) = q(u, i)^\gamma \cdot \left(1 + \beta_{c(i)} r(i) + \eta_{c(i)} g(i) + \mu m(i) + \nu v(i)\right) \cdot (1 + \rho p(h)) - \delta C(i) - \lambda D(i, L). \quad (1)$$

The exponent $\gamma \geq 0$ modulates the influence of base relevance. Contract-aware weights $\beta_{c(i)}$ and $\eta_{c(i)}$ determine how profitability $r(i)$ and plan-gap $g(i)$ are traded off under each settlement type. The parameters μ and ν control the magnitude of the contract-value signal and premium flag, respectively. The pacing term is scaled by ρ . The penalties for cap proximity and similarity are controlled by δ and λ . Table 4 summarizes the tunable parameters of Equation (1).

Table 4. Tunable parameters of the PMOR scoring function (Equation (1)).

Parameter	Role
$\gamma \geq 0$	exponent modulating base relevance
$\beta_{c(i)}$	contract-specific weight on profitability $r(i)$
$\eta_{c(i)}$	contract-specific weight on plan-gap $g(i)$
μ	weight on contract-value signal $m(i)$
ν	weight on premium flag $v(i)$
ρ	scaling of pacing term $p(h)$
δ	cap-penalty coefficient
λ	diversity-penalty coefficient

Equation (1) is designed to be interpretable: each module corresponds to a business concept and the multiplicative structure provides a clear pathway for how relevance interacts with business modifiers.

3.3.2. Contract-Aware Weights

The settlement model affects what constitutes an optimal allocation of exposure. Under pay-per-action contracts, the score reflects the economic structure of completed actions while operating within agreed cap constraints, so that exposure remains consistent with the contractually expected volume. Under fixed-fee contracts, the objective shifts toward meeting agreed delivery plans while avoiding excessive lead volume beyond what is economically rational.

We implement this distinction through contract-specific weights:

$$(\beta_{\text{fixed}}, \eta_{\text{fixed}}) = (0.2, 0.7), \quad (\beta_{\text{cpa}}, \eta_{\text{cpa}}) = (0.6, 0.3). \quad (2)$$

These values reflect a stronger emphasis on plan-gap under fixed-fee contracts and a stronger emphasis on profitability under pay-per-action contracts. The weights are not treated as immutable constants; rather, they serve as a plausible starting point that can be tuned by segment or market region if necessary.

3.3.3. Objective Formulation and Selection

The re-ranking policy trades off a small set of objectives that were chosen to match the operational levers available at inference time. For a slate L_u , we consider:

$$\begin{aligned} & \max \underbrace{\sum_{i \in L_u} q(u, i)}_{\text{relevance}}, \quad \max \underbrace{\sum_{i \in L_u} r(i)}_{\text{lead value}}, \quad \max \underbrace{\sum_{i \in L_u} g(i)}_{\text{plan fulfillment}}, \\ & \min \underbrace{\sum_{i \in L_u} C(i)}_{\text{cap pressure}}, \quad \min \underbrace{\sum_{i \in L_u} D(i, L)}_{\text{redundancy}}. \end{aligned} \quad (3)$$

Relevance, lead value and plan fulfillment are maximized; cap pressure and within-slate redundancy are minimized. These objectives were selected because each corresponds to a directly controllable business or user-experience lever (user–item match, monetizable lead quality, contractual pacing, exposure safety and slate diversity), rather than to a

latent quantity that would require a learned surrogate. PMOR combines them by explicit parametric scalarization at ranking time, which keeps each term interpretable and lets operators re-weight strategy without retraining the base model.

The objectives are not statistically independent: for example, high-value inventory that is behind its plan tends to score well on both lead value and plan fulfillment, while aggressive exposure of a single investment raises both cap pressure and redundancy. Because the terms are partially correlated and can conflict, the scalarization weights and the guardrail caps are what prevent any single objective from dominating the slate; Section 4 reports the component-wise contributions that make this balance explicit.

3.4. Feature Construction and Guardrails

3.4.1. Base Relevance Score Calibration

Pipeline Step 1 (Figure 1): the calibrated base relevance $q(u, i)$ from which the item score is built.

The base model score is treated as a proxy for conversion probability. In production, such a score may be a similarity in embedding space or a model output that is not naturally calibrated to $[0, 1]$. For the offline evaluation, a calibrated score $q(u, i) \in [0, 1]$ is derived from the logged base output via two steps:

$$q_{\text{raw}}(u, i) \leftarrow \begin{cases} 0, & \text{if the raw score is a sentinel value,} \\ q_{\text{raw}}(u, i), & \text{otherwise,} \end{cases} \quad (4)$$

followed by clipping:

$$q(u, i) = \text{clip}(q_{\text{raw}}(u, i), 0, 1). \quad (5)$$

While simple, this guardrail is crucial for stability. A missing relevance score should not become artificially dominant through exponentiation, nor should it produce undefined behavior.

3.4.2. Plan-Gap and Profitability Signals

Pipeline Step 2 (Figure 1): the multiplicative business adjustments applied on top of base relevance (continued in the contract-value and contract-aware-weight subsections).

The plan-gap signal $g(i)$ represents how far an investment is from its target delivery. It is normalized to $[0, 1]$, where values near 1 indicate under-delivery and values near 0 indicate that the plan is on track or completed. In the offline logs used here, $g(i)$ is available as a precomputed component and is clipped to $[0, 1]$ to maintain consistent scaling.

The profitability signal $r(i)$ is a normalized indicator of lead value. In production systems, this signal is typically derived from lead prices, margins, or a contract-aware transformation thereof. For the offline evaluation, $r(i)$ is treated as a stable input feature already normalized to $[0, 1]$, again applying clipping when needed.

3.4.3. Contract-Value Signal and Premium Flag

The contract-value signal $m(i) \in [-1, 1]$ is a bounded, symmetric modifier of the score derived from a normalized monthly budget proxy, allowing controlled adjustments in either direction around a neutral reference point. The signal is derived from a monthly budget proxy available in daily snapshots. Let $B(i) \geq 0$ be the raw budget proxy for the investment associated with i . We apply a log transform for scale stabilization:

$$b_{\log}(i) = \log(1 + B(i)). \quad (6)$$

A min–max normalization yields $\tilde{b}(i) \in [0, 1]$ within a chosen segmentation, with a fallback to 0.5 when the segment is degenerate:

$$\tilde{b}(i) = \begin{cases} \frac{b_{\log}(i) - b_{\min}}{b_{\max} - b_{\min}}, & \text{if } b_{\max} > b_{\min}, \\ 0.5, & \text{if } b_{\max} = b_{\min}. \end{cases} \quad (7)$$

Finally, rescale to $[-1, 1]$:

$$m(i) = 2\tilde{b}(i) - 1. \quad (8)$$

This construction ensures that the impact of $\mu m(i)$ is symmetric and bounded, preventing the budget signal from becoming numerically unstable when raw budgets vary over orders of magnitude.

The premium flag $v(i) \in \{0, 1\}$ marks investments in the upper tail of an effective lead-value distribution. The flag is computed from an effective lead price $\pi(i)$ defined in Section 4.1. Within each contract type, a percentile threshold (90th percentile in the evaluated configuration) defines premium items:

$$v(i) = \mathbf{1}\{\pi(i) \geq P_{90}(\pi)\}. \quad (9)$$

3.4.4. Cap Ratio and Cap Penalty

Pipeline Step 3 (Figure 1): the first of the two subtractive guardrail penalties (cap proximity; within-slate diversity follows).

Cap progress is represented by a ratio $\text{cap_ratio}(\text{inv}(i))$ that measures predicted or realized delivery at the investment level relative to the plan. When direct forecasted realization is available in snapshots, it is used. Otherwise, a fallback based on leads-to-date and lead limit is applied:

$$\text{cap_ratio}(\text{inv}(i)) = \begin{cases} \text{estimated_perc_realization}(\text{inv}(i)), & \text{if available,} \\ \frac{\text{leads}(\text{inv}(i))}{\max(\text{lead_limit}(\text{inv}(i)), 1)}, & \text{otherwise.} \end{cases} \quad (10)$$

A cap penalty $C(i) \in [0, 1]$ is then defined with a threshold at 0.7, increasing linearly to 1 at full saturation:

$$C(i) = \begin{cases} 0, & \text{if } \text{cap_ratio}(\text{inv}(i)) < 0.7, \\ \frac{\text{cap_ratio}(\text{inv}(i)) - 0.7}{0.3}, & \text{if } 0.7 \leq \text{cap_ratio}(\text{inv}(i)) < 1.0, \\ 1, & \text{if } \text{cap_ratio}(\text{inv}(i)) \geq 1.0. \end{cases} \quad (11)$$

The threshold at 0.7 is a pragmatic compromise: it avoids premature penalization when an investment is comfortably below plan, yet it activates early enough to prevent late-day surges that exceed caps.

3.4.5. Diversity Penalty for Within-Slate Redundancy

Diversity is implemented via a similarity function defined on discrete identifiers:

$$\text{sim}(i, j) = \max\left(w_{\text{inv}}\mathbf{1}\{\text{inv}(i) = \text{inv}(j)\}, w_{\text{dev}}\mathbf{1}\{\text{dev}(i) = \text{dev}(j)\}, w_{\text{city}}\mathbf{1}\{\text{city}(i) = \text{city}(j)\}\right), \quad (12)$$

with weights satisfying $w_{\text{inv}} \geq w_{\text{dev}} \geq w_{\text{city}} \geq 0$. The diversity penalty relative to the partial slate L is:

$$D(i, L) = \max_{j \in L} \text{sim}(i, j). \quad (13)$$

The max operator captures a “dominant similarity” interpretation. If two properties belong to the same investment, that redundancy is treated as more severe than merely sharing a developer or city. In the evaluated configuration, the weights are:

$$w_{\text{inv}} = 1.0, \quad w_{\text{dev}} = 0.5, \quad w_{\text{city}} = 0.25. \quad (14)$$

3.4.6. Pacing Signal and the Role of ρ

The pacing term is intended to regulate the delivery rate over the day. In a full production implementation, a dynamic pacing signal compares planned and served exposure up to hour h :

$$p(h) = \text{clip}\left(\frac{T_h - S_h}{T_h + \varepsilon}, -1, 1\right), \quad (15)$$

where T_h denotes cumulative exposure target by time h , S_h is cumulative served exposure and ε is a small stability constant. The multiplier $(1 + \rho p(h))$ therefore lies in $[1 - \rho, 1 + \rho]$.

In the offline evaluation, full served exposure trajectories were not reconstructed. The analysis therefore uses a fixed $p(h) = p_h$ to reflect a configuration where pacing is enabled but not dynamically driven by hour-level state. This choice allows the score decomposition to isolate pacing as a bounded multiplicative factor while acknowledging that production pacing is inherently stateful.

3.5. Sequential Slate Construction for PMOR

Pipeline Step 5 (Figure 1): the greedy sequential slate construction that turns per-item scores into the final TOP-K list.

Because the score includes a term $D(i, L)$ that depends on the partial slate, the final ranking is constructed sequentially. For each request, candidates are iteratively selected by recomputing $S(u, i)$ with respect to the current slate L .

The sequential procedure is computationally feasible because K is small and candidate sets are modest in size after first-stage retrieval. With average candidate count about 44, the per-request complexity scales as $\mathcal{O}(K|\mathcal{C}_u|)$, with low constant factors because score components are simple arithmetic and discrete comparisons.

3.6. Reference Rankers: Baseline and LPR

This subsection defines part of the evaluation setup used in Section 4 rather than a component of the PMOR method: the reference rankers against which PMOR is compared. The offline profit proxy used as the evaluation objective is now defined together with the research questions in that section.

The evaluation compares three ranking strategies. The baseline strategy ranks candidates solely by calibrated base relevance $q(u, i)$, using a deterministic tie-breaker on item identifiers when necessary. This strategy reflects a purely model-driven ranking without explicit business objectives.

The second comparator is a legacy production re-ranker (LPR). LPR corresponds to the ranking used in the marketplace at the time of data collection. In the offline logs, LPR is represented by a “final” score that includes existing business adjustments. LPR applies business-aware adjustments on top of the base relevance score—incorporating signals such as contract state, inventory availability and legacy commercial adjustments—to produce the “final” score recorded in logs. Its internals are proprietary to the marketplace operator and were not modified during data collection, so we treat LPR as a black-box baseline, comparing PMOR against the legacy score as observed. For evaluation, LPR ranks candidates by this legacy score.

The third strategy is PMOR, which constructs TOP-K using Algorithm 1 with the scoring function in Equation (1).

3.7. Parameter Selection and Optimization

The score in Equation (1) contains parameters $\Theta = \{\gamma, \mu, \nu, \rho, \delta, \lambda\}$ and potentially a fixed p_h when pacing is not dynamic. In practice, parameter selection must respect both business objectives and user experience constraints, such as not degrading relevance quality beyond an acceptable threshold. The offline study focuses on profit proxy as the primary objective while using bounded features and penalties to avoid extreme behavior.

The evaluated configuration uses a set of parameters obtained through an offline search procedure that samples candidate parameter vectors within bounded ranges and selects those maximizing $\hat{R}_{\text{PMOR}}$ on a train split, with a validation split to reduce overfitting. The search is based on random sampling augmented with local refinement around top-performing candidates. This approach is motivated by two factors. First, the objective is non-convex and depends on sequential selection. Second, the relative scaling of score components is not known a priori, so grid search becomes inefficient in higher dimensions.

Because the objective depends on sequential slate construction, computational efficiency is important. Candidate features are cached per request and the sequential procedure is applied only to the small number of candidates in TOP-K selection.

Sensitivity to the scalarization weights is asymmetric across parameters. The weights on the business-value and cap-penalty terms have the largest effect on the profit proxy, because they directly govern how much billable, plan-behind inventory is promoted and how hard the layer brakes near contractual caps. The diversity weight has a comparatively small effect on the proxy but bounds worst-case slate redundancy. This means the layer can be tuned primarily along the value/cap axis, while the diversity term acts as a low-cost guardrail rather than a driver of the reported gains.

3.8. Deterministic Component-Wise Contribution Analysis

A practical advantage of Equation (1) is that it enables exact attribution of the final score to individual components. Rather than using surrogate explainability methods, one can decompose the score into additive contributions that correspond directly to business modules. For a fixed $p(h)$, define:

$$S_{\text{core}}(u, i) = q(u, i)^\gamma \cdot (1 + \rho p(h)). \quad (16)$$

The multiplicative business term can be expanded into additive components scaled by $q(u, i)^\gamma$:

$$\begin{aligned} S_r(u, i) &= q(u, i)^\gamma \cdot \beta_{c(i)} r(i) \cdot (1 + \rho p(h)), \\ S_g(u, i) &= q(u, i)^\gamma \cdot \eta_{c(i)} g(i) \cdot (1 + \rho p(h)), \\ S_m(u, i) &= q(u, i)^\gamma \cdot \mu m(i) \cdot (1 + \rho p(h)), \\ S_v(u, i) &= q(u, i)^\gamma \cdot \nu v(i) \cdot (1 + \rho p(h)). \end{aligned} \quad (17)$$

The penalty components are:

$$S_{\text{cap}}(i) = \delta C(i), \quad S_{\text{div}}(i, L) = \lambda D(i, L). \quad (18)$$

The full score is therefore:

$$S(u, i) = S_{\text{core}}(u, i) + S_r(u, i) + S_g(u, i) + S_m(u, i) + S_v(u, i) - S_{\text{cap}}(i) - S_{\text{div}}(i, L). \quad (19)$$

This decomposition supports two complementary analyses. First, one can compute the distribution of each component over selected items in TOP-K to understand which drivers dominate. Second, one can perform counterfactual ablations by removing a component

(e.g., setting $\delta = 0$ or $m(i) = 0$) and recomputing scores, which indicates how much each module contributes on average.

4. Results

The offline evaluation is organized around four research questions. RQ1: Does contract-aware re-ranking improve the business profit proxy relative to a model-only baseline and to the legacy production re-ranker? RQ2: How is any gain distributed across rank positions and settlement models? RQ3: How much does each score component contribute to the final ranking, and which guardrails are active? RQ4: How stable is the re-ranked slate relative to the base ranking? Each question is answered by a dedicated part of this section: RQ1 by the overall profit-proxy comparison (Section 4.2) and the per-request pairwise test (Section 4.3); RQ2 by the breakdowns across rank positions and settlement models (Sections 4.4 and 4.5), together with the cap and diversity exposure outcomes in Sections 4.7 and 4.8; RQ3 by the component-wise contribution results (Section 4.10); and RQ4 by the ranking-stability and list-overlap analysis (Section 4.6). The evaluation mechanism that these results rely on—the offline profit proxy—is specified first, and Section 4.11 closes with the threats to validity.

4.1. Profit Proxy for Offline Evaluation

This subsection specifies the evaluation mechanism used to answer the research questions above: an offline expected-revenue proxy that is consistently computable from production logs, together with its effective-price, conversion and position-discount components.

In production, true revenue depends on downstream user actions, attribution windows and billing rules. Such signals are typically not available at the granularity of offline logs. We therefore evaluate using an expected-revenue proxy that is consistently computable offline and that reflects two key facts. The proxy reflects two assumptions: item value varies across the candidate set, and user attention decreases with rank position.

4.1.1. Effective Lead Price

To compare across settlement models, we define an effective lead price $\pi(i)$:

$$\pi(i) = \begin{cases} \text{lead_price}(i), & c(i) = \text{cpa}, \\ \text{estimated_lead_price}(i), & c(i) = \text{fixed}. \end{cases} \quad (20)$$

When `estimated_lead_price` is missing under fixed-fee contracts, a fallback uses budget divided by lead limit or by forecasted leads:

$$\pi(i) \leftarrow \pi(i) \text{ fillna}\left(\frac{B(i)}{\max(\text{lead_limit}(i), 1)}\right) \text{ fillna}\left(\frac{B(i)}{\max(\text{leads_estimation}(i), 1)}\right). \quad (21)$$

4.1.2. Conversion Probability Proxy and Billability

Because $q(u, i) \in [0, 1]$ by construction (Equation (5)), we use the calibrated score directly as a conversion probability proxy:

$$p(i) = q(u, i). \quad (22)$$

The proxy explicitly uses the base relevance score rather than the legacy production score to avoid over-counting business adjustments. This “apples-to-apples” choice ensures that differences in profit proxy arise from ranking changes and item selection, not from substituting a business-adjusted score as a probability.

Billability is modeled as a binary guardrail derived from cap ratio:

$$b(i) = \mathbf{1}\{\text{cap_ratio}(\text{inv}(i)) < 1.0\}. \quad (23)$$

4.1.3. Position Discounting and Expected-Revenue Proxy

Let w_j denote position weights. We use a DCG-style discount:

$$w_j = \frac{1}{\log_2(j+1)}, \quad j \in \{1, \dots, K\}. \quad (24)$$

For a ranking strategy A producing a slate $i_A(1), \dots, i_A(K)$, the per-request proxy is:

$$\hat{R}_A(u) = \sum_{j=1}^K \pi(i_A(j)) \cdot p(i_A(j)) \cdot w_j \cdot b(i_A(j)), \quad (25)$$

and the dataset-level proxy aggregates over requests:

$$\hat{R}_A = \sum_{u \in \mathcal{U}} \hat{R}_A(u). \quad (26)$$

Table 5 summarizes the components of the per-request profit proxy (Equation (25)).

Table 5. Components of the per-request expected-revenue proxy (Equation (25)).

Component	Role
$\pi(i)$	effective lead price (Equations (20) and (21))
$p(i)$	conversion probability proxy = $q(u, i)$ (Equation (22); distinct from pacing signal $p(h)$)
w_j	DCG position weight (Equation (24))
$b(i)$	billability indicator (Equation (23))

4.2. Overall Profit Proxy Comparison

This subsection addresses RQ1: whether contract-aware re-ranking improves the profit proxy relative to the baseline and LPR.

Table 6 reports the total expected-revenue proxy $\hat{R}$ for TOP-K = 3. The table also reports the share of billable positions and the fraction of selected positions with missing effective lead price after applying the fallback rules in Equation (21). A weighted average conversion proxy $\bar{q}_w$ is computed as the ratio of $\sum \hat{R}$ to the sum of value-position weights without probability, providing a sanity check that ranking changes do not trivially arise from selecting systematically lower-probability items.

To preserve commercial confidentiality, all proxy values in this section are reported as indices where the baseline total equals 100.

Table 6. Overall expected-revenue proxy (indexed, baseline = 100) for TOP-K = 3 under three ranking strategies.

Ranker	$\sum \hat{R}$	Lift vs. Baseline	$\sum$ Unit Value-Pos	$\bar{q}_w$	Billable	Missing Price
baseline	100.0	0.0%	100.0	69.9%	32.30%	3.36%
LPR	328.4	228.4%	321.7	71.3%	37.85%	0.32%
PMOR	487.4	387.4%	504.2	67.5%	78.22%	1.16%

The overall lift from baseline to LPR indicates that the legacy production logic already shifts exposure toward higher-value and more billable inventory compared with pure

relevance sorting. PMOR further increases the profit proxy by a large margin over LPR. The billable share difference is particularly pronounced, suggesting that cap-aware behavior is a dominant driver of the proxy objective. We stress that these figures, including the 387.4% lift over the model-only baseline, are computed against the offline profit proxy rather than realized revenue; they should therefore be read as proxy-level improvements that may overstate live business impact until confirmed in an online test.

4.3. Per-Request Pairwise Comparison: LPR vs. PMOR

Aggregate results can be sensitive to outliers. To assess whether the advantage of PMOR is broad-based, per-request proxy sums are compared. For each request, let $\Delta(u) = \hat{R}_{LPR}(u) - \hat{R}_{PMOR}(u)$. Table 7 summarizes the distribution of $\Delta(u)$.

Table 7. Per-request comparison of LPR and PMOR on 4219 requests.

Metric	Value
Requests compared	4219
LPR higher (%)	8.94%
PMOR higher (%)	54.21%
Tie (%)	36.86%
Mean $\Delta(u)$	−32.32
Median $\Delta(u)$	−5.83
Sum of $\Delta(u)$	−159.1

The comparison suggests that PMOR outperforms LPR for a majority of requests rather than relying on a small number of extreme gains. The high tie rate is consistent with the billability gate in Equation (23): when both rankers select mostly non-billable items for a request, the proxy becomes less discriminative.

4.4. Breakdown by Rank Position

This subsection and the next address the first half of RQ2: how the gain is distributed across rank positions (the settlement-model split follows).

Position discounting implies that improvements at rank 1 are more valuable than improvements at rank 3. Table 8 shows the contribution of each position to the overall proxy for each ranker. The parentheses report the share of total proxy contributed by that position.

Table 8. Expected-revenue proxy (indexed, baseline = 100) by rank position for TOP-K = 3.

Rank	Baseline	LPR	PMOR
1	49.7	179.2	286.9
2	28.3	89.2	125.2
3	22.0	59.9	75.3

PMOR is more top-heavy, concentrating a larger fraction of value at rank 1. This behavior is consistent with the sequential selection policy combined with cap and diversity penalties: the method tends to prioritize a high-quality, billable and non-redundant first item, which receives the highest position weight.

4.5. Breakdown by Settlement Model

To evaluate whether PMOR behaves consistently across settlement models, proxy results are reported separately for pay-per-action (CPA) and fixed-fee contracts. Table 9 reports proxy sums by contract type, together with the share of rows of each contract type in TOP-3 and the lift versus baseline within the same contract segment.

Table 9. Expected-revenue proxy (indexed) by settlement model (CPA vs. fixed-fee) for TOP-K = 3.

Contract Type	Baseline	LPR	PMOR
CPA	60.8 (49.19%)	206.4 (48.54%; 239.6%)	300.7 (54.05%; 394.7%)
fixed-fee	39.2 (50.81%)	121.9 (51.46%; 210.9%)	186.7 (45.95%; 376.3%)

The improvement of PMOR over LPR appears in both settlement types, indicating that the method does not merely optimize for one contract model at the expense of the other. The contract-aware weights in Equation (2) provide a plausible mechanism: profitability is emphasized more in CPA, while plan-gap receives higher weight in fixed-fee agreements.

4.6. Ranking Stability and List Overlap

This subsection addresses RQ4: how stable the re-ranked slate is relative to the base ranking.

A business-aware re-ranker may substantially adjust slates relative to relevance-only ranking, which can affect user experience and complicate deployment. To quantify how much the slate changes, Table 10 reports overlap metrics computed between rankers, including average overlap in TOP-3 and agreement on TOP-1.

Table 10. Slate overlap metrics between rankers (TOP-3 overlap and TOP-1 agreement).

Pair	Avg. TOP-3 Overlap	Full TOP-3 Match	TOP-1 Agreement
LPR vs. PMOR	35.22%	5.71%	31.17%
baseline vs. LPR	23.95%	3.58%	12.59%
baseline vs. PMOR	16.32%	2.39%	6.57%

The overlap indicates that PMOR is not a minor adjustment. This magnitude of change is aligned with its purpose as a business-aware layer, but it motivates guardrails in deployment, such as monitoring relevance metrics and user engagement and introducing ramp-up strategies.

Concretely, PMOR reorders the base list by promoting candidates whose business adjustments (billable, high-value or plan-behind) outweigh a modest relevance gap, and by demoting candidates that trigger cap or diversity penalties. The stability and overlap statistics above quantify this: the re-ranked TOP-K retains much of the base relevance mass while redirecting a controlled share of exposure toward billable inventory, which is the intended behavior of the layer.

4.7. Cap Distribution and Billable Exposure

Since the profit proxy gates non-billable items, understanding cap distribution is critical. Table 11 shows the distribution of selected TOP-3 positions across three cap regions for PMOR: below threshold (no cap penalty), linear region (penalty increasing) and saturated (cap ratio at least 1.0).

Table 11. Distribution of cap regions among PMOR-selected TOP-3 positions.

Cap Region	Count	Share
below ($\text{cap_ratio} < 0.7$)	8241	65.43%
linear ($0.7 \leq \text{cap_ratio} < 1.0$)	1598	12.69%
saturated ($\text{cap_ratio} \geq 1.0$)	2756	21.88%

The sum of below and linear regions implies that 78.12% of PMOR-selected positions are below cap, consistent with the billable share in Table 6. Residual saturated selections remain nontrivial, suggesting either limited availability of billable candidates for some requests or conservative penalty weights relative to other score components.

For additional context, the offline summaries indicate that baseline and LPR select substantially more saturated positions. In the evaluated sample, the share of positions with $\text{cap_ratio} \geq 1.0$ is approximately 67.70% for baseline, 62.15% for LPR and 21.88% for PMOR. This shift is consistent with PMOR's explicit cap penalty in Equation (1).

4.8. Diversity Outcomes and Redundancy Reduction

A re-ranking layer that focuses on commercial alignment may inadvertently increase redundancy by selecting multiple items from a high-value investment. PMOR addresses this risk through the within-slate diversity penalty. In the evaluated configuration, the penalty activates rarely but strongly when it does.

In the offline sample, the fraction of requests with duplicate investment identifiers in TOP-3 is approximately 12.3% for baseline, 5.8% for LPR and 0.5% for PMOR. This reduction aligns with the structure of Equation (13), where exact investment matches carry the highest similarity weight.

The activation reasons are also informative. When the diversity penalty activates, it is driven by investment duplication rather than developer- or city-level duplication. This observation suggests that in real-estate recommendation, redundancy is primarily a within-investment issue and that city-level redundancy may be less problematic or already controlled by other constraints.

4.9. Most Affected Investment Groups

Operationally, large ranking changes tend to concentrate on a limited set of investments where caps, lead value and plan status interact. To support auditing, we identify the most affected investment groups Table 12 by absolute difference in aggregated profit proxy between LPR and PMOR. To preserve anonymity, investment identifiers are replaced by arbitrary group labels.

Table 12. Top-10 anonymized investment groups with the largest absolute difference in expected-revenue proxy (indexed) between LPR and PMOR.

Group	$\hat{R}_{\text{LPR}}$	$\hat{R}_{\text{PMOR}}$	$\Delta (\text{LPR} - \text{PMOR})$
G01	1.60	21.44	−19.84
G02	2.60	17.55	−14.96
G03	0.11	7.91	−7.80
G04	23.76	30.26	−6.50
G05	6.70	0.78	5.92
G06	0.00	5.07	−5.07
G07	5.98	0.38	5.61
G08	2.21	7.00	−4.79
G09	0.00	4.53	−4.53
G10	3.33	0.73	2.60

The table highlights that differences can be asymmetric. Some groups benefit dramatically under PMOR, likely due to being billable and high value but under-exposed under LPR. Other groups lose exposure, plausibly because they were near or beyond caps or because they were redundant relative to other selected items. This concentration suggests that deployment validation can focus on a manageable subset of investment groups to diagnose whether changes reflect intended policy or data quality issues in cap and value features.

4.10. Component-Wise Contribution Results

This subsection addresses RQ3: how much each score component contributes to the final ranking and which guardrails are active.

This section reports the component-wise contribution analysis for PMOR-selected TOP-3 positions. The analysis focuses on the additive terms in Equation (19). For each selected item, each component contribution is computed and summary statistics are aggregated over all selected items.

4.10.1. Global Contribution Magnitudes

Table 13 reports mean contribution, mean absolute contribution, median absolute contribution and the fraction of cases where the contribution is positive or negative. The “share” column expresses each component’s fraction of the total mean absolute contribution across the components shown.

Table 13. Global component contribution statistics for PMOR-selected TOP-3 items.

Component	Mean	Mean Abs	Median Abs	% > 0	% < 0	Share (Mean Abs)
<i>m</i> (contract value)	0.936	1.129	1.061	92.96%	7.04%	22.58%
core	0.779	0.779	0.795	76.55%	0.00%	15.57%
<i>g</i> (plan-gap)	0.416	0.416	0.401	74.56%	0.00%	8.31%
cap penalty	−0.397	0.442	0.000	0.00%	33.61%	8.85%
<i>r</i> (profitability)	0.264	0.285	0.184	71.52%	0.00%	5.69%
<i>v</i> (premium)	0.108	0.111	0.000	18.11%	0.00%	2.22%
diversity penalty	−0.003	0.005	0.000	0.00%	0.17%	0.10%

Two patterns stand out. The contract-value signal *m* dominates contribution magnitude, suggesting that the evaluated parameterization places substantial weight on normalized budget. The cap penalty is sparse but impactful: its median is zero because it is inactive for many items, yet its mean absolute contribution is comparable to plan-gap. The diversity penalty is extremely sparse, consistent with its role as a guardrail that activates only when redundancy occurs.

4.10.2. Contributions by Rank Position

Table 14 shows mean absolute contributions by rank position. The diversity penalty is absent at rank 1 by construction because there is no previously selected item, while cap penalties increase in lower ranks, consistent with penalized items being pushed down the slate.

Table 14. Mean absolute component contributions by rank position for PMOR-selected TOP-3 items.

Component	Rank 1	Rank 2	Rank 3
<i>m</i>	1.208 (24.15%)	1.121 (22.42%)	1.072 (21.44%)
core	0.780 (15.61%)	0.776 (15.50%)	0.791 (15.81%)
<i>g</i>	0.335 (6.71%)	0.355 (7.09%)	0.317 (6.34%)
cap penalty	0.221 (4.42%)	0.397 (7.94%)	0.593 (11.86%)
<i>r</i>	0.286 (5.72%)	0.254 (5.09%)	0.246 (4.91%)
<i>v</i>	0.139 (2.79%)	0.087 (1.74%)	0.093 (1.85%)
diversity penalty	0.000 (0.00%)	0.003 (0.05%)	0.005 (0.11%)

4.10.3. Contributions by Settlement Model

A core design claim is that PMOR behaves contract-aware through $\beta_{c(i)}$ and $\eta_{c(i)}$. Table 15 reports mean absolute contributions separately for CPA and fixed-fee items.

The profitability component r is substantially larger for CPA than for fixed-fee, while plan-gap g is larger for fixed-fee than for CPA. This behavior matches the intent of the contract-aware weights in Equation (2). The cap penalty is also stronger for CPA, consistent with CPA contracts exhibiting sharper billability boundaries and more frequent saturation in the operational state.

Table 15. Mean absolute component contributions by settlement model for PMOR-selected TOP-3 items.

Component	CPA	Fixed-Fee
m	1.274 (25.47%)	0.961 (19.22%)
core	0.815 (16.29%)	0.741 (14.82%)
g	0.254 (5.09%)	0.419 (8.38%)
cap penalty	0.463 (9.25%)	0.306 (6.13%)
r	0.411 (8.22%)	0.113 (2.27%)
v	0.156 (3.12%)	0.047 (0.94%)
diversity penalty	0.005 (0.10%)	0.004 (0.07%)

4.10.4. Counterfactual Ablations of Score Components

To assess component importance in a causal sense, we compute counterfactual changes in the score when one component is removed. This analysis answers the question of how the score would change if, for instance, the contract-value signal were disabled or cap penalties were removed. Table 16 reports the mean change Δ and the distribution of absolute changes.

Table 16. Counterfactual score changes when disabling individual components in PMOR (evaluated on selected TOP-3 items).

Counterfactual	Δ Mean	Mean Abs	Median Abs	% $\Delta > 0$	% $\Delta < 0$	% $\Delta = 0$
Disable m	−0.873	1.081	1.004	0.85%	75.70%	23.45%
Disable g	−0.315	0.315	0.322	0.02%	74.54%	25.44%
Disable r	−0.226	0.226	0.232	0.00%	71.52%	28.48%
Disable cap penalty	0.396	0.397	0.000	33.61%	0.00%	66.39%
Disable v	−0.104	0.104	0.000	0.00%	18.11%	81.89%
Disable pacing	−0.081	0.081	0.083	0.00%	76.55%	23.45%
Disable diversity penalty	0.003	0.003	0.000	0.17%	0.00%	99.83%

Disabling m produces the largest average score decrease, consistent with Table 13. Removing the cap penalty increases the score when the penalty was active, reflecting that the cap module is a deliberate regulator of exposure. The diversity penalty affects almost none of the items, but its effect can be substantial when redundancy would otherwise occur.

4.11. Threats to Validity

We close the validation with the threats to validity that qualify the answers to RQ1–RQ4, together with the actions taken to mitigate them.

Construct validity. The evaluation objective is an offline expected-revenue proxy (Section 4.1) rather than realized revenue, so the answers to RQ1 and RQ2 are only as valid as the proxy’s lead-value, billability and position-discount assumptions. To mitigate this, the proxy uses the calibrated base relevance rather than the business-adjusted score as its conversion term (Equation (22)), so that measured differences arise from ranking and item selection rather than from re-using a business score as a probability, and every component is reported separately (Table 5) so its influence can be audited.

Internal validity. Because the logs are observational, unobserved confounders in logged feedback could influence the component-wise attribution behind RQ3. As-of joins against daily snapshots mitigate look-ahead bias, and the contribution analysis is exact and deterministic (Section 4.10) rather than a learned surrogate, so the reported magnitudes are reproducible directly from the scoring rule. The per-request pairwise test (Section 4.3) further checks that the RQ1 gain is broad-based rather than driven by a few extreme requests.

External validity. Results are drawn from a single anonymized marketplace and reporting window, so they may not transfer to other inventory mixes, settlement structures or interfaces. The settlement-model breakdown (Section 4.5) shows the RQ2 effect holds across both contract types in this sample, but broader transfer remains open. These threats jointly motivate the online validation discussed under Limitations in the Discussion.

5. Discussion

The three contributions stated in the Introduction are addressed as follows. The parameterized contract-aware scoring function is defined in Section 3.1, specified by the scoring rule in Equation (1) and the slate-level objective formulation in Equation (3); the inference-time sequential slate construction is given by the greedy procedure in Algorithm 1; and the deterministic component-wise contribution analysis is developed and reported in the results, where each module's effect on the ranking is quantified and ablated.

The empirical results suggest that a structured, contract-aware re-ranking layer can substantially improve business objectives without requiring retraining of the underlying relevance model. The results also show that business-aware ranking can substantially reshape recommendation lists, which motivates careful deployment.

5.1. Why PMOR Improves the Profit Proxy

The dominant driver of profit proxy improvement appears to be increased billable exposure. The proxy metric in Equation (25) is explicitly zeroed when items are not billable. Under this metric, a ranker that successfully avoids saturated investments will improve $\hat{R}$ even if it selects slightly lower conversion probabilities, as long as the selected items remain sufficiently relevant.

PMOR enforces this behavior through two mechanisms. The cap penalty in Equation (11) begins to apply before full saturation, discouraging selections as an investment approaches its cap. The billability indicator in Equation (23) is not used directly by the re-ranker, but it shapes the evaluation metric; consequently, cap-aware behavior yields large proxy improvements. The component-wise analysis reinforces this interpretation: cap contributions are sparse but increase with lower ranks, indicating that penalized items are pushed down the slate.

A secondary driver is the contract-value signal $m(i)$. The global contribution analysis shows that this component has the largest contribution magnitude. This pattern should not be interpreted as commercial signals overriding relevance. Rather, it reflects the evaluated parameterization and the fact that $m(i)$ is normalized to $[-1, 1]$ and applies broadly across items. When $m(i)$ is active and large, it can become the principal determinant of exposure allocation, which may be desirable from a contractual perspective but requires careful scrutiny for user experience.

5.2. Trade-Offs and Potential Risks

The observed gains are encouraging, but they come with trade-offs. First, the overlap analysis shows that PMOR can differ substantially from baseline and from LPR. Large changes can improve business objectives while risking relevance degradation if constraints

are misspecified or if business signals are noisy. This risk is amplified when the base relevance exponent γ is small, as $q(u, i)^\gamma$ compresses differences and allows business terms to take precedence. This behavior may be appropriate when business constraints are strict, but it should be monitored with user-centric metrics.

Second, the use of a budget-derived $m(i)$ can lead to exposure concentration if not balanced by cap and diversity controls. Although $m(i)$ is bounded, its weight μ can induce systematic preferences. Whether such preferences are acceptable depends on contractual commitments and marketplace design. In some environments, allocating exposure in proportion to contract value is expected; in others, it may not align with broader marketplace fairness norms or may reduce long-term supply diversity.

Third, diversity control in PMOR is intentionally conservative. The diversity penalty activates rarely in the evaluated sample, primarily on investment-level duplication. This behavior may be appropriate if within-investment redundancy is the principal user-experience issue. If developer-level or geographic diversity is a goal, additional constraints or a revised similarity definition would be required.

Although the diversity penalty accounts for only a small share of the total component contribution, we retain it deliberately. Its purpose is not to drive the profit proxy but to bound a tail risk: without it, sequential selection can produce near-duplicate slates (multiple listings from the same investment), which degrade user experience even when average business metrics look favorable. As a low-cost guardrail that rarely binds but prevents pathological slates, it is kept in its current form.

5.3. Marketplace Fairness and Exposure Bias

Because PMOR promotes billable and high-value inventory, exposure can become correlated with advertiser budget. The exposure caps and pacing signal act as a partial fairness guardrail by bounding how much attention any single investment or developer can capture, preventing over-concentration on the largest advertisers. However, the layer does not currently enforce an explicit supplier-fairness constraint, and a budget-correlated exposure bias may remain. Provider-side fairness objectives and minimum-exposure guarantees for smaller developers are a natural extension and are noted as future work.

5.4. Estimating User Satisfaction

The present evaluation optimizes a business profit proxy and does not directly measure user satisfaction. In deployment, satisfaction would be estimated from behavioral proxies already logged by the platform—click-through rate on recommended slates, dwell time, contact/lead submission and downstream conversion—and compared against the model-only baseline in an online test. Incorporating such user-centric metrics alongside the business proxy is important to ensure that commercial gains do not degrade the user experience, and is part of the planned online validation.

5.5. Operational Considerations for Deployment

A key benefit of the proposed approach is that the re-ranking policy is parameterized and can be updated without retraining the base model. This is useful when contract terms, budgets, or business strategy change frequently. However, operational deployment requires attention to several practical aspects.

Latency constraints may necessitate efficient feature retrieval and caching. Although the scoring function is lightweight, it depends on up-to-date cap ratios and sometimes on dynamic pacing signals. These signals must be computed reliably and made available at inference time.

Monitoring should include both business metrics (billable share, plan fulfillment, revenue proxies) and user-centric metrics (click-through rates, downstream engagement,

satisfaction proxies). Because the offline objective focuses on profit proxy, it does not guarantee preservation of user experience. A staged rollout with controlled experimentation is therefore appropriate.

5.6. Limitations of the Offline Proxy Objective

The expected-revenue proxy is a useful offline tool but has limitations. The effective lead price $\pi(i)$ is a proxy rather than an observed marginal revenue. In fixed-fee contracts, the as-if lead price depends on budget and estimated leads, which may be noisy. The conversion proxy $p(i)$ is derived from base relevance and may not be calibrated. Additionally, the proxy assumes independence across positions and does not model user browsing behavior beyond simple position discounting.

These limitations do not invalidate the comparison, because the proxy is applied consistently across rankers. However, they imply that offline improvements may not translate linearly into online revenue. The proxy is best interpreted as an indicator that the ranking policy shifts exposure toward more valuable, billable inventory under the assumed model.

Several limitations bound the interpretation of the reported gains. First, the headline improvement is measured against a constructed profit proxy rather than realized revenue; the proxy is a modeling choice, and the magnitude of the lift should be read as a proxy-level result that may overstate real-world impact until confirmed online. Second, we report no online A/B validation of click-through, lead generation or user satisfaction; establishing that the offline gains translate to live outcomes is the primary item of future work. Third, we do not benchmark against population-based multi-objective optimizers such as NSGA-II or explicit Pareto-front methods: PMOR is a deterministic, inference-time policy designed for controllability and exact attribution rather than frontier search, so such methods are complementary baselines we plan to compare against in future work. Fourth, the aggregate proxy figures are reported as point estimates without dispersion measures; bootstrap confidence intervals and per-request variance are a straightforward addition for a future, computation-inclusive revision.

6. Conclusions

This paper presented a contract-aware, multi-objective re-ranking layer for online real-estate recommendation. The proposed method (PMOR) combines calibrated relevance with explicit business drivers and operational guardrails, including contract-type-aware weighting, exposure caps and within-slate diversity control. The deterministic structure of the scoring rule enables exact component-wise contribution analysis and counterfactual ablations, supporting transparent tuning and auditing.

Offline evaluation on production logs from an anonymized marketplace showed that PMOR substantially improves a business-oriented proxy metric compared with both a model-only baseline and a legacy production re-ranker (LPR). The primary driver is improved billable exposure under cap constraints, with additional gains associated with contract-value and plan-gap signals.

Framed against the four research questions, the study answers them as follows. RQ1: contract-aware re-ranking improves the offline profit proxy—PMOR reaches an indexed proxy of 487.4, a 387.4% lift over the model-only baseline and a 48.4% lift over LPR. RQ2: the gain is concentrated at the top of the slate (PMOR is markedly more top-heavy at rank 1) and holds across both settlement models, with the TOP-3 billable share rising to 78.22% from 37.85% under LPR. RQ3: the component-wise analysis identifies the contract-value signal and the cap penalty as the dominant drivers, while the diversity penalty acts as a rarely-binding but strong guardrail. RQ4: the re-ranked slate is a substantial but

controlled adjustment of the base ranking (TOP-3 overlap of 35.22% and TOP-1 agreement of 31.17% against LPR), consistent with its role as a business-aware layer. As noted in the threats to validity, these are proxy-level results and RQ1–RQ2 in particular require online confirmation.

Consistent with this, the most important next step is to confirm the RQ1–RQ2 proxy gains in a live online experiment. Future work includes validating the approach in on-line experiments, incorporating richer user interaction models, learning segment-specific parameters under guardrails and exploring alternative diversity definitions that capture geographic novelty and developer-level variety without compromising contractual pacing.

Author Contributions: Conceptualization, M.B. and J.K.; Methodology, B.A., M.B., B.K., A.K. and J.K.; Software, B.A., M.B., B.K., A.K. and J.K.; Validation, P.W., B.A., M.B., B.K., A.K. and J.K.; Formal Analysis, B.A., M.B. and J.K.; Investigation, B.A., M.B. and J.K.; Resources, B.A., M.B., B.K., A.K. and J.K.; Data Curation, B.A., M.B., B.K. and A.K.; Writing—Original Draft Preparation, B.A. and J.K.; Writing—Review and Editing, B.A., M.B., B.K., A.K., M.K., S.S., P.W. and J.K.; Supervision, P.W., M.B. and J.K.; Project Administration, P.W., M.B. and J.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Data Availability Statement: The datasets analyzed in this study originate from an anonymized industry partner and contain commercially sensitive operational signals. Aggregated statistics and methodology details are provided in the article. Additional information may be made available by the corresponding author upon reasonable request, subject to contractual and privacy constraints. The scoring rule, algorithm and evaluation protocol are fully specified in the article, so the method can be reproduced on comparable marketplace logs even though the proprietary dataset cannot be released.

Conflicts of Interest: Authors Bogdan Arct, Mateusz Bieniek, Bartłomiej Kanabus, Aleksander Kozłowski and Piotr Wetmański were employed by the company Property Group sp. z o.o. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

CPA	Cost per action (pay-per-action settlement)
DCG	Discounted Cumulative Gain
LPR	Legacy production re-ranker
MMR	Maximal Marginal Relevance
NDCG	Normalized Discounted Cumulative Gain
P90	90th percentile
PMOR	Proposed multi-objective re-ranker
TOP-K	The top <i>K</i> ranked items returned by the system

References

1. Abdollahpouri, H.; Burke, R.; Mobasher, B. Recommender Systems as Multistakeholder Environments. In *UMAP'17: Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*; ACM: New York, NY, USA, 2017; pp. 347–348. [[CrossRef](#)]
2. Choi, H.; Mela, C.F. Monetizing online marketplaces. *Mark. Sci.* **2019**, *38*, 948–972. [[CrossRef](#)]
3. Biswas, A.; Patro, G.K.; Ganguly, N.; Gummadi, K.P.; Chakraborty, A. Toward Fair Recommendation in Two-sided Platforms. *ACM Trans. Web* **2022**, *16*, 1–34. [[CrossRef](#)]
4. Acharya, A.; Gao, S.; Saha, A.; Ocejó, B.; Basu, K.; Selvaraj, K.; Mazumder, R.; Gupta, A.; Agrawal, P. Optimizing for Member Value in an Edge Building Marketplace. In *CIKM'23: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*; ACM: New York, NY, USA, 2023; pp. 5–14. [[CrossRef](#)]

5. Bonaldo, G.G.; Nogueira, P.F.; Torovets, T.; da Silva, T.F. Beyond Relevance: A Demand Balancer Model for Rental Platforms with Single-Unit Inventory. In *WSDM'25: Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining*; ACM: New York, NY, USA, 2025; pp. 1086–1087. [\[CrossRef\]](#)
6. Lee, P.C.; Liu, P.J. Digital information-seeking of homebuyers on real estate platforms in Taiwan. *Electron. Libr.* **2025**, *43*, 342–367. [\[CrossRef\]](#)
7. Gharahighehi, A.; Pliakos, K.; Vens, C. Recommender systems in the real estate market—A survey. *Appl. Sci.* **2021**, *11*, 7502. [\[CrossRef\]](#)
8. Muñoz, E.G.; Parraga-Alava, J.; Meza, J.; Proaño Morales, J.J.; Ventura, S. Housing fuzzy recommender system: A systematic literature review. *Heliyon* **2024**, *10*, e26444. [\[CrossRef\]](#) [\[PubMed\]](#)
9. Henríquez-Miranda, C.; Ríos-Pérez, J.; Sanchez-Torres, G. Recommender systems in real estate: A systematic review. *Bull. Electr. Eng. Inform.* **2025**, *14*, 2156–2170. [\[CrossRef\]](#)
10. Venkatesh, C.; Oberoi, H.; Goyal, A.; Sikka, N. RE-RecSys: An End-to-End system for recommending properties in Real-Estate domain. In *CODS-COMAD'24: Proceedings of the 7th Joint International Conference on Data Science & Management of Data (11th ACM IKDD CODS and 29th COMAD)*; ACM: New York, NY, USA, 2024; pp. 558–562. [\[CrossRef\]](#)
11. Ricci, F.; Rokach, L.; Shapira, B. *Recommender Systems Handbook*; Springer: Boston, MA, USA, 2010; pp. 1–35. [\[CrossRef\]](#)
12. Covington, P.; Adams, J.; Sargin, E. Deep Neural Networks for YouTube Recommendations. In *RecSys'16: Proceedings of the 10th ACM Conference on Recommender Systems*; ACM: New York, NY, USA, 2016; pp. 191–198. [\[CrossRef\]](#)
13. Gomez-Urbe, C.A.; Hunt, N. The Netflix Recommender System: Algorithms, Business Value, and Innovation. *ACM Trans. Manag. Inf. Syst. (TMIS)* **2016**, *6*, 1–19. [\[CrossRef\]](#)
14. Higley, K.; Oldridge, E.; Ak, R.; Rabhi, S.; De Souza Pereira Moreira, G. Building and Deploying a Multi-Stage Recommender System with Merlin. In *RecSys'22: Proceedings of the 16th ACM Conference on Recommender Systems*; ACM: New York, NY, USA, 2022; pp. 632–635. [\[CrossRef\]](#)
15. Liu, W.; Qin, J.; Tang, R.; Chen, B. Neural Re-ranking for Multi-stage Recommender Systems. In *RecSys '22: Proceedings of the 16th ACM Conference on Recommender Systems*; ACM: New York, NY, USA, 2022; pp. 698–699. [\[CrossRef\]](#)
16. Zhao, B.; Qi, H.; Xu, G.; Ma, M.; Zhao, X.; Mei, F.; Xu, S.; Hu, J. A Hybrid Cross-Stage Coordination Pre-ranking Model for Online Recommendation Systems. In *WWW'25: Companion Proceedings of the ACM on Web Conference 2025*; ACM: New York, NY, USA, 2025; pp. 621–630. [\[CrossRef\]](#)
17. Qian, H.; Wu, Q.; Zhang, K.; Zhang, Z.; Gu, L.; Zeng, X.; Zhou, J.; Gu, J. Scope-aware re-ranking with gated attention in feed. In *WSDM '22: Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*; ACM: New York, NY, USA, 2022; pp. 804–812. [\[CrossRef\]](#)
18. Xi, Y.; Lin, J.; Liu, W.; Dai, X.; Zhang, W.; Zhang, R.; Tang, R.; Yu, Y. A Bird's-eye View of Reranking: From List Level to Page Level. In *WSDM'23: Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*; ACM: New York, NY, USA, 2023; pp. 1075–1083. [\[CrossRef\]](#)
19. Zeng, A.; Yu, H.; He, H.; Ni, Y.; Li, Y.; Zhou, J.; Miao, C. Enhancing E-commerce Recommender System Adaptability with Online Deep Controllable Learning-To-Rank. In *Proceedings of the AAAI Conference on Artificial Intelligence*; AAAI: Washington, DC, USA, 2021; Volume 35, pp. 15214–15222. [\[CrossRef\]](#)
20. Liu, Q.; Zheng, K.; Huang, R.; Li, W.; Cai, K.; Chai, Y.; Niu, Y.; Hui, Y.; Han, B.; Mou, N.; et al. RecFlow: An Industrial Full Flow Recommendation Dataset. In *International Conference on Learning Representations 2025 (ICLR 2025)*; Curran Associates, Inc.: Red Hook, NY, USA, 2025; pp. 90319–90336.
21. Joachims, T.; Swaminathan, A.; de Rijke, M. Deep Learning with Logged Bandit Feedback. In *Proceedings of the International Conference on Learning Representations*; Curran Associates, Inc.: Red Hook, NY, USA, 2018.
22. Sankara Moorthy, K.; Fallah, M.H.; Govindaraj, M.; Rama Sree, S.; Kanmani Jebaseeli, S.; Ugli, J.T.M. Neural Re-Ranking Models in Research Article Recommendations. *Indian J. Inf. Sources Serv.* **2025**, *15*, 330–338. [\[CrossRef\]](#)
23. Kayaalp, E.; Keskinöz, M. MH-NAIS: A Multi-Head Attention Extension of the Neural Attentive Item Similarity Model. In *2025 9th International Symposium on Innovative Approaches in Smart Technologies (ISAS)*; IEEE: Piscataway, NJ, USA, 2025. [\[CrossRef\]](#)
24. Kim, M.; Kim, W. Task-Oriented Collaborative Graph Embedding Using Explicit High-Order Proximity for Recommendation. *Big Data Res.* **2023**, *33*, 100382. [\[CrossRef\]](#)
25. Chen, H.; He, J.; Xu, W.; Feng, T.; Liu, M.; Song, T.; Yao, R.; Qiao, Y. Enhanced Multi-Relationships Integration Graph Convolutional Network for Inferring Substitutable and Complementary Items. In *Proceedings of the AAAI Conference on Artificial Intelligence*; AAAI: Washington, DC, USA, 2023; Volume 37, pp. 4157–4165. [\[CrossRef\]](#)
26. Chen, J.; Yang, X.; Yang, C.; Bao, J.; Guo, Z.; Li, Y.; Shi, C. CORONA: A Coarse-to-Fine Framework for Graph-based Recommendation with Large Language Models. In *SIGIR'25: Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*; ACM: New York, NY, USA, 2025; pp. 2048–2058. [\[CrossRef\]](#)

27. Gao, H.; Huang, W.; Wu, F.; Zhao, J.; Zheng, C.; Liu, H. LLM Enhancers for GNNs: An Analysis from the Perspective of Causal Mechanism Identification. In *Proceedings of the 42nd International Conference on Machine Learning*; JMLR: New York, NY, USA, 2025; Volume 267, pp. 18585–18611.
28. Nerrise, F.; Huang, E.W.; Ji, X.; Subbian, K.; Koutra, D. GraFS: An Integrated GNN-LLM Approach for Inferring Best Functional Substitute Products. In *CIKM'25: Proceedings of the 34th ACM International Conference on Information and Knowledge Management*; ACM: New York, NY, USA, 2025; pp. 5047–5051. [\[CrossRef\]](#)
29. Noor, K.; Fatima, U. Meta Learning Strategies for Comparative and Efficient Adaptation to Financial Datasets. *IEEE Access* **2025**, *13*, 24158–24170. [\[CrossRef\]](#)
30. Burges, C.; Shaked, T.; Renshaw, E.; Lazier, A.; Deeds, M.; Hamilton, N.; Hullender, G. Learning to Rank using Gradient Descent. In *ICML'05: Proceedings of the 22nd International Conference on Machine Learning*; ACM: New York, NY, USA, 2005; pp. 89–96. [\[CrossRef\]](#)
31. Cao, Z.; Qin, T.; Liu, T.Y.; Tsai, M.F.; Li, H. Learning to Rank: From Pairwise Approach to Listwise Approach. In *ICML'07: Proceedings of the 24th International Conference on Machine Learning*; ACM: New York, NY, USA, 2007; Volume 227, pp. 129–136. [\[CrossRef\]](#)
32. Marler, R.; Arora, J. Survey of Multi-Objective Optimization Methods for Engineering. *Struct. Multidiscip. Optim.* **2004**, *26*, 369–395. [\[CrossRef\]](#)
33. Mesran, M.; Sihombing, P.; Efendi, S.; Zarlis, M. SAGOMECON: An Adaptive Constraint-Based Optimization Method for Multi-Criteria Decision-Making in Collaborative Industrial Networks. *J. Appl. Data Sci.* **2025**, *6*, 2682–2695. [\[CrossRef\]](#)
34. Shukla, A.; Sathyanarayana, J.; Banerjee, D. Sample-Rank: Weak Multi-Objective Recommendations Using Rejection Sampling. In *2020 International Conference on Data Mining Workshops (ICDMW)*; IEEE: Piscataway, NJ, USA, 2020; pp. 683–689. [\[CrossRef\]](#)
35. Agah, S.; Schaeffer, S.; Peifer, M.; Sharma, N.; Maheshwari, A.; Hamidian, S. Pareto-Optimal Solution: Optimizing Engagement and Revenue. In *RecSys'25: Proceedings of the Nineteenth ACM Conference on Recommender Systems*; ACM: New York, NY, USA, 2025; pp. 1034–1037. [\[CrossRef\]](#)
36. Chen, S.; Wang, Y.; Wen, Z.; Li, Z.; Zhang, C.; Zhang, X.; Lin, Q.; Zhu, C.; Xu, J. Controllable Multi-Objective Re-ranking with Policy Hypernetworks. In *KDD'23: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*; ACM: New York, NY, USA, 2023; pp. 3855–3864. [\[CrossRef\]](#)
37. Cunha, T.; Partalas, I.; Nguyen, P. Juggler: Multi-stakeholder ranking with meta-learning. In *MORS@ RecSys*; ACM: New York, NY, USA, 2021; Volume 2959.
38. Ie, E.; Jain, V.; Wang, J.; Narvekar, S.; Agarwal, R.; Wu, R.; Cheng, H.T.; Chandra, T.; Boutilier, C. SlateQ: A Tractable Decomposition for Reinforcement Learning with Recommendation Sets. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*; AAAI: Washington, DC, USA, 2019; pp. 2592–2599. [\[CrossRef\]](#) [\[PubMed\]](#)
39. Carbonell, J.; Goldstein, J. The Use of MMR, Diversity-Based Reranking for Reordering Documents and Producing Summaries. In *SIGIR'98: Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*; ACM: New York, NY, USA, 1999. [\[CrossRef\]](#)
40. Ziegler, C.-N.; McNee, S.M.; Konstan, J.A.; Lausen, G. Improving recommendation lists through topic diversification. In *WWW '05: Proceedings of the 14th International Conference on World Wide Web*; ACM: New York, NY, USA, 2005. [\[CrossRef\]](#)
41. Kulesza, A.; Taskar, B. Determinantal Point Processes for Machine Learning. *Found. Trends® Mach. Learn.* **2012**, *5*, 123–286. [\[CrossRef\]](#)
42. Singh, A.; Joachims, T. Fairness of Exposure in Rankings. In *KDD'18: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*; ACM: New York, NY, USA, 2018; pp. 2219–2228. [\[CrossRef\]](#)
43. Biega, A.; Gummadi, K.; Weikum, G. Equity of Attention: Amortizing Individual Fairness in Rankings. In *SIGIR'18: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*; ACM: New York, NY, USA, 2018. [\[CrossRef\]](#)
44. Chifu, A.G.; Mothe, J.; Ullah, M.Z. Fair exposure of documents in information retrieval: A community detection approach. In *Circle 2020*; ACM: New York, NY, USA, 2020; Volume 2621.
45. Sonboli, N.; Eskandarian, F.; Burke, R.; Liu, W.; Mobasher, B. Opportunistic Multi-aspect Fairness through Personalized Re-ranking. In *UMAP'20: Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*; ACM: New York, NY, USA, 2020; pp. 239–247. [\[CrossRef\]](#)
46. Marras, M.; Boratto, L.; Ramos, G.; Fenu, G. Regulating Group Exposure for Item Providers in Recommendation. In *SIGIR'22: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*; ACM: New York, NY, USA, 2022; pp. 1839–1843. [\[CrossRef\]](#)
47. Xu, C.; Chen, S.; Xu, J.; Shen, W.; Zhang, X.; Wang, G.; Dong, Z. P-MMF: Provider Max-min Fairness Re-ranking in Recommender System. In *WWW'23: Proceedings of the ACM Web Conference 2023*; ACM: New York, NY, USA, 2023; pp. 3701–3711. [\[CrossRef\]](#)
48. Naghiaei, M.; Rahmani, H.A.; Deldjoo, Y. CPFair: Personalized Consumer and Producer Fairness Re-ranking for Recommender Systems. In *SIGIR'22: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*; ACM: New York, NY, USA, 2022; pp. 770–779. [\[CrossRef\]](#)

49. Wu, G.; Feng, Z.; Chen, S.; Wu, H.; Xue, X.; Xiao, J.; Fan, G.; Chen, H.; Li, J. FairSort: Learning to Fair Rank for Personalized Recommendations in Two-Sided Platforms. *IEEE Trans. Knowl. Data Eng.* **2025**, *37*, 641–654. [\[CrossRef\]](#)
50. Raj, A.; Ekstrand, M.D. Towards Optimizing Ranking in Grid-Layout for Provider-Side Fairness. In *Advances in Information Retrieval; Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*; Springer: Cham, Switzerland, 2024; Volume 14612, pp. 90–105. [\[CrossRef\]](#)
51. Xu, C.; Xu, J.; Ding, Y.; Zhang, X.; Qi, Q. FairSync: Ensuring Amortized Group Exposure in Distributed Recommendation Retrieval. In *WWW'24: Proceedings of the ACM Web Conference 2024*; ACM: New York, NY, USA, 2024; pp. 1092–1102. [\[CrossRef\]](#)
52. Mansoury, M.; Abdollahpouri, H.; Pechenizkiy, M.; Mobasher, B.; Burke, R. A Graph-Based Approach for Mitigating Multi-Sided Exposure Bias in Recommender Systems. *ACM Trans. Inf. Syst.* **2022**, *40*, 1–31. [\[CrossRef\]](#)
53. Pisani, A. On the Longitudinal Impact of Exposure Bias in Recommender Systems. In *Advances in Information Retrieval; Lecture Notes in Computer Science*; Springer: Cham, Switzerland, 2025; Volume 15576, pp. 178–183. [\[CrossRef\]](#)
54. Zhernov, A.; Dvijotham, K.D.; Lobov, I.; Calian, D.A.; Gong, M.; Chandrashekar, N.; Mann, T.A. The NodeHopper: Enabling Low Latency Ranking with Constraints via a Fast Dual Solver. In *KDD'20: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*; ACM: New York, NY, USA, 2020; pp. 1285–1294. [\[CrossRef\]](#)
55. Wang, W.; Chen, C.; Yu, C.; Hao, X.; Jin, J.; Zhang, W.; Wang, Y.; Hao, J.; Wang, J.; Li, H.; et al. Learning adaptive display exposure for real-time advertising. In *CIKM'19: Proceedings of the 28th ACM International Conference on Information and Knowledge Management*; ACM: New York, NY, USA, 2019; pp. 2595–2603. [\[CrossRef\]](#)
56. Balseiro, S.; Gur, Y. Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium. *Manag. Sci.* **2019**, *65*, 3952–3968. [\[CrossRef\]](#)
57. Geyik, S.C.; Chowdhury, L.; Raudies, F.; Pu, W.; Shen, J. Impression Pacing for Jobs Marketplace at LinkedIn. In *CIKM'20: Proceedings of the 29th ACM International Conference on Information & Knowledge Management*; ACM: New York, NY, USA, 2020; pp. 2445–2452. [\[CrossRef\]](#)
58. Celli, A.; Colini-Baldeschi, R.; Kroer, C.; Sodomka, E. The Parity Ray Regularizer for Pacing in Auction Markets. In *WWW'22: Proceedings of the ACM Web Conference 2022*; ACM: New York, NY, USA, 2022; pp. 162–172. [\[CrossRef\]](#)
59. Dai, L.; Lyu, K.; Zhang, C.; Zhao, G.; Zu, Z.; Wang, L.; Zheng, B. Percentile Risk-Constrained Budget Pacing for Guaranteed Display Advertising in Online Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*; AAAI: Washington, DC, USA, 2024; Volume 38, pp. 7987–7994. [\[CrossRef\]](#)
60. Chen, Q.; Nguyen, P.H.; Gligorijevic, D. Optimization-Based Budget Pacing in eBay Sponsored Search. In *WWW'24: Companion Proceedings of the ACM Web Conference 2024*; ACM: New York, NY, USA, 2024; pp. 328–337. [\[CrossRef\]](#)
61. Stram, R.; Abboud, R.; Shtoff, A.; Somekh, O.; Raviv, A.; Koren, Y. Mystique: A Budget Pacing System for Performance Optimization in Online Advertising. In *WWW'24: Companion Proceedings of the ACM Web Conference 2024*; ACM: New York, NY, USA, 2024; pp. 433–442. [\[CrossRef\]](#)
62. Dündar, U.; Abou Nabout, N.; Skiera, B. Getting the pace right: Performance of budget allocation heuristics in online advertising. *Int. J. Res. Mark.* **2025**, *in press*. [\[CrossRef\]](#)
63. Hajiaghayi, M.; Springer, M. Analysis of a Learning Based Algorithm for Budget Pacing. In *AAMAS'23: Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*; ACM: New York, NY, USA, 2023; pp. 2709–2711.
64. Laarabi, M.; Chegri, B.; Mohammadia, A.M.; Lafriouni, K. Smart Contracts Applications in Real Estate: A Systematic Mapping Study. In *2022 2nd International Conference on Innovative Research in Applied Science, Engineering and Technology (IRASET)*; IEEE: Piscataway, NJ, USA, 2022. [\[CrossRef\]](#)
65. Shehu, A.S.; Pinto, A.; Correia, M.E. A Decentralised Real Estate Transfer Verification based on Self-Sovereign Identity and Smart Contracts. In *Proceedings of the 19th International Conference on Security and Cryptography SECRYPT*; SciTePress: Setúbal, Portugal, 2022; Volume 1, pp. 469–476. [\[CrossRef\]](#)
66. Preethi, Ulla, M.M.; Sapna, R.; Devadas, R.M.; Hiremani, V.; Pavithra, N. Automated Escrow Services in Real Estate: Enhancing Trust and Efficiency with Smart Contracts. In *Intelligent Sustainable Systems; Lecture Notes in Networks and Systems*; Springer: Cham, Switzerland, 2026; Volume 1729, pp. 226–235. [\[CrossRef\]](#)
67. Sowmya, G.; Alankruthi, K.; Harika, J.; Nagini, T.; Nagalakshmi, M.; Kumar, J.P. Smart Contracts for Real Estate Sales-A New Era of Efficiency and Transparency. In *2024 IEEE 6th International Conference on Cybernetics, Cognition and Machine Learning Applications (ICCCMLA)*; IEEE: Piscataway, NJ, USA, 2024; pp. 228–231. [\[CrossRef\]](#)
68. Redekar, S.; Bhagwat, S.; Upasani, K.; Jha, S.; Thorat, V. Real Estate Management System using Blockchain. In *2024 3rd International Conference on Applied Artificial Intelligence and Computing (ICAAIC)*; IEEE: Piscataway, NJ, USA, 2024; pp. 1455–1460. [\[CrossRef\]](#)
69. Miriam, J.; Rajashree; Nafeesmuneera, M.; Saranya, V.; Srihas, G.; Kumar, B.G. Decentralization of real estate using blockchain. *AIP Conf. Proc.* **2025**, *3257*, 020096. [\[CrossRef\]](#)
70. Lalitha, R.; Sankari, M.; Sree, H.; Gunda, G.V. A decentralized real estate system using blockchain technology. *AIP Conf. Proc.* **2025**, *3257*, 020098. [\[CrossRef\]](#)

71. Ujjwal; Gupta, P.; Peter, J.S.P. Intelligent Housing System with Cryptographically Secure Ownership and Dynamic Price Bidding. In *2023 6th International Conference on Recent Trends in Advance Computing (ICRTAC)*; IEEE: Piscataway, NJ, USA, 2023; pp. 469–473. [\[CrossRef\]](#)
72. Amin, M.R.M.; Abbas, M.S.; Ismail, A.F.; Rahmat, R.; Zolkefle, N.Z.N.; Nordin, M.N. Blockchain Systems and Applications: Evaluating the Opportunities of the Blockchain Technology in Various Industries. In *2023 3rd International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE)*; IEEE: Piscataway, NJ, USA, 2023; pp. 1146–1149. [\[CrossRef\]](#)
73. Mottaghi, S.F.; Steininger, B.I.; Yanagawa, N. Real Estate Insights: The current state and the new future of tokenization in real estate. *J. Prop. Investig. Financ.* **2024**, *42*, 614–620. [\[CrossRef\]](#)
74. Mistrangelo, P.; Tagliabue, L.C.; Tezel, A. Digital Environment Definition for Property Tokenization Uptake in Italy. In *EC3 Conference 2023*; European Council on Computing in Construction: Sint-Niklaas, Belgium, 2023. [\[CrossRef\]](#)
75. Mandloi, J.; Bansal, P. Smart Contract Security and Privacy Taxonomy, Tools, and Challenges. In *Data Science and Security*; Lecture Notes in Networks and Systems; Springer: Singapore, 2021; Volume 290, pp. 159–167. [\[CrossRef\]](#)
76. Dhillon, D.; Diksha; Mehrotra, D. Smart Contract Vulnerabilities: Exploring the Technical and Economic Aspects. In *Blockchain Transformations*; Signals and Communication Technology; Springer: Cham, Switzerland, 2024; pp. 81–91. [\[CrossRef\]](#)
77. Meloche, R.; Sivakumar, D.; Anda, A.A.; Alfuhaid, S.; Amyot, D.; Logrippo, L.; Mylopoulos, J. A Web-Based Environment for the Specification and Generation of Smart Legal Contracts. In *Compliance for Artificial Intelligence Systems*; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2026; Volume 14377, pp. 11–23. [\[CrossRef\]](#)
78. Dolat-abadi, H.M.; Rasi, M.; Sadat, S.A. *Smart Contracts as a Third Party Coordinator: Tools for Implementing Agreements in E-Business Management*; IGI Global: Hershey, PA, USA, 2023; pp. 123–151. [\[CrossRef\]](#)
79. Lin, S. Multi-objective Optimization of Recommendation System Based on Graph Attention Network. In *Proceedings of the 2nd International Conference on Educational Knowledge and Informatization (EKI'24)*; ACM: New York, NY, USA, 2024; pp. 367–373. [\[CrossRef\]](#)

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.